



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE

United States Patent and Trademark Office

Address: COMMISSIONER FOR PATENTS

P.O. Box 1450

Alexandria, Virginia 22313-1450

www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/767,182	01/29/2004	Pazhani Pillai	SUN03-03(P9403)	4649
51344 7590 10/15/2010 BROOKS KUSHMAN P.C. /Oracle America/ SUN / STK 1000 TOWN CENTER, TWENTY-SECOND FLOOR SOUTHFIELD, MI 48075-1238				
EXAMINER				
HOANG, HIEU T				
ART UNIT		PAPER NUMBER		
2452				
MAIL DATE		DELIVERY MODE		
10/15/2010		PAPER		

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Office Action Summary

Application No.

10/767,182

Applicant(s)

PILLAI ET AL.

Examiner

HIEU HOANG

Art Unit

2452

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --
Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

- 1) ☒ Responsive to communication(s) filed on 25 August 2010.
- 2a) ☐ This action is **FINAL**. 2b) ☒ This action is non-final.
- 3) ☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

- 4) ☒ Claim(s) 1-27 is/are pending in the application.
- 4a) Of the above claim(s) _____ is/are withdrawn from consideration.
- 5) ☐ Claim(s) _____ is/are allowed.
- 6) ☒ Claim(s) 1-27 is/are rejected.
- 7) ☐ Claim(s) _____ is/are objected to.
- 8) ☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

- 9) ☐ The specification is objected to by the Examiner.
- 10) ☐ The drawing(s) filed on _____ is/are: a) ☐ accepted or b) ☐ objected to by the Examiner.
Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).
Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).
- 11) ☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

- 12) ☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
- a) ☐ All b) ☐ Some * c) ☐ None of:
1. ☐ Certified copies of the priority documents have been received.
 2. ☐ Certified copies of the priority documents have been received in Application No. _____.
 3. ☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

- 1) ☐ Notice of References Cited (PTO-892)
- 2) ☐ Notice of Draftperson's Patent Drawing Review (PTO-948)
- 3) ☐ Information Disclosure Statement(s) (PTO/SB/CD)
Paper No(s)/Mail Date _____
- 4) ☐ Interview Summary (PTO-413)
Paper No(s)/Mail Date _____
- 5) ☐ Notice of Informal Patent Application
- 6) ☐ Other: _____

DETAILED ACTION

1. A request for continued examination under 37 CFR 1.114, including the fee set forth in 37 CFR 1.17(e), was filed in this application after final rejection. Since this application is eligible for continued examination under 37 CFR 1.114, and the fee set forth in 37 CFR 1.17(e) has been timely paid, the finality of the previous Office action has been withdrawn pursuant to 37 CFR 1.114. Applicant's submission filed on 08/25/2010 has been entered.
2. Claims 1-27 are pending.

Response to Arguments

3. Applicant's arguments have been fully considered but they are not persuasive. Applicant argues that the prior art does not teach wherein the other data access tasks are of the type including a first request, a first response, at least one subsidiary request, and at least one subsidiary response, and wherein context resources used to service the other data access tasks are pre-emptible. It is maintained that Saha does teach the limitation ([0058-0059], resource is pre-emptible as soon as the partial task is finished transferring; the resource is available for other tasks such as other partial tasks of the same task).

Claim Objections

4. Claims 1-27 are objected to because of the following informalities: independent claims recite in the newly amended limitation "a first request, a first response, at least

one subsidiary request, at least one subsidiary response" should refer back to the first request, the first response, the at least one subsidiary request, and the at least one subsidiary response previously recited. Appropriate correction is required.

Claim Rejections - 35 USC § 101

5. 35 U.S.C. 101 reads as follows:

Whoever invents or discovers any new and useful process, machine, manufacture, or composition of matter, or any new and useful improvement thereof, may obtain a patent therefor, subject to the conditions and requirements of this title.

6. Claim 27 is rejected under 35 USC 101 because the claimed invention is directed to nonstatutory subject matter. The claim is drawn to a "computer program product having a computer-readable medium" comprising stored data. The specification is silent regarding the definitions of the terms "computer program product" and "computer-readable medium". Thus, applying the broadest reasonable interpretation in light of the specification and taking into account the meaning of the words in their ordinary usage as they would be understood by one of ordinary skill in the art, the claim as a whole covers a transitory signal, which does not fall within the definition of a process, a machine, manufacture, or composition of matter (*In re Nuijten*).

Amendments may be made to narrow the claim to cover only statutory embodiments to avoid a rejection under 35 U.S.C. § 101 by adding the limitation "non-transitory" to the claim. *Cf Animals - Patentability*, 1077 Off. Gaz. Pat. Office 24 (April 21, 1987) (suggesting that applicants add the limitation "non-human" to a claim covering a multicellular organism to avoid a rejection under 35 U.S.C. § 101). Such an

amendment would typically not raise the issue of new matter, even when the specification is silent because the broadest reasonable interpretation relies on the ordinary and customary meaning that includes non-transitory embodiments and transitory signals per se. *Subject Matter Eligibility of Computer Readable Media*, 1351 OG 212 (February 23, 2010).

Claim Rejections - 35 USC § 103

7. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

8. **Claims 1-11, 13-23, and 25-27 are rejected under 35 U.S.C. 103(a) as being unpatentable over Saha et al. (US 2004/0117375, hereafter Saha), in view of Boyd et al. (US 2004/0049580, hereafter Boyd).**

9. For claim 1, Saha discloses a method for performing remote access commands between nodes, the method comprising:

issuing a first request from a first node to a second node, the first request requesting a data access be performed between the first node and the second node ([0060], fig. 5, client writes data to server's database);

receiving, at the first node, a first response from the second node that partially completes the data access ([0060], fig. 5, acknowledgement of write completion associated with a buffer segment at the server);

issuing at least one subsidiary request from the first node to the second node to further complete the data access between the first node and the second node, the at least one subsidiary request based on an amount of partial completion of the data access between the first node and the second node ([0060], fig. 5, after receiving an ACK indicating a buffer read completion from the server, the client can refill its buffer and initiate another subsequent write request); and

receiving, from the second node in response to the at least one subsidiary request, at least one corresponding subsidiary response that further completes the data access between the first node and the second node (fig. 5, another ACK for completion of the subsequent read request);

wherein context resources in a channel adapter of the first node used to service the data access task are pre-emptible after only partial completion of the data access task ((fig. 2, 5, [0059-0061], memory resource preempted for a data chunk when the data chunk has been read or written and acknowledged, the data chunk is the partial task) between the first response and a first subsidiary request for use to support other data access tasks by the channel adapter performing a context switch, and context resources in the channel adapter used to service the data access task are pre-emptible between a subsidiary response and a subsequent subsidiary request for use to support other data access tasks by the channel adapter performing a context switch (fig. 2, 5,

[0059-0061], memory resource preempted for a data chunk between an ACK response and transferring of the next data block; the client software for emptying its buffers and sending ACKs for context switching is read as the channel adapter).

wherein the other data access tasks are of the type including a first request, a first response, at least one subsidiary request, and at least one subsidiary response, and wherein context resources used to service the other data access tasks are pre-emptible ([0058-0059], resource is pre-emptible as soon as the partial task is finished transferring; the resource is available for other tasks such as other partial tasks of the same task).

Saha does not explicitly disclose a data access task.

However, Boyd discloses a data access task (fig. 4, a work request is a task comprising a plurality of data segments or frames)

Therefore, it would have been obvious for one skilled in the art at the time of the invention to combine the teachings of Saha and Boyd to apply the tasking scheme of Boyd to the DMA buffering management scheme of Saha to monitor or manage completion of task segments.

10. Claims 26 and 27 are rejected for the same rationale given in claim 1.

11. For claim 25, Saha discloses a method for performing read remote access commands between nodes ([0022]), the method comprising:

issuing a first read remote access command from a first node to a second node (fig. 5, [0061], read RDMA);

receiving at least one first response from the second node that partially completes a data access associated with the first remote access command ([0061], receiving response with data amount in buffer size);

pre-empting context resources associated with issuance of the first remote access command from the first node prior to completion of the data access associated with the first remote access command (fig. 1, 5, [0059]-[0061], client emptying its allocated buffer blocks after finishing reading the received data);

issuing a second remote access command from the first node using the pre-empted context resources ([0061], client preparing an ACK to the server for each available emptied client buffer); and

issuing, from the first node to the second node, a series of subsidiary remote access commands derived from the first remote access command to receive corresponding subsidiary responses from the second node to complete the data access associated with the first read remote access command ([0061], client sending one or more ACKs to the server indicating more available client buffers so that more read data can be sent to the client).

Saha does not explicitly disclose a data access task.

However, Boyd discloses a data access task (fig. 4, a work request is a task comprising a plurality of data segments or frames)

Therefore, it would have been obvious for one skilled in the art at the time of the invention to combine the teachings of Saha and Boyd to apply the tasking scheme of Boyd to the DMA buffering management scheme of Saha to monitor or manage completion of task segments.

12. For claim 2, Saha-Boyd discloses the invention as in claim 1. Saha-Boyd further discloses issuing the first request from the first node comprises:

detecting an application request in a request queue, the application request identifying the data access task to be performed between the first and second node (RDMA read/write request); and wherein the method comprises: repeating issuing at least one subsidiary request and receiving at least one corresponding subsidiary response between the first and second nodes until the data access task is totally complete between the first and second nodes (Saha, [0063], queue segments for subsequent buffer read/write iteration).

13. For claim 3, Saha-Boyd discloses the invention as in claim 2. Saha-Boyd further discloses allocating context resources in the first node for receipt of the first response and for receipt of the at least one subsidiary response ([0061], client buffers for receiving read data responses), the context resources allocated to support receipt of the first response and for receipt of the at least one subsidiary response that contain data in an amount not exceeding a data allotment to support at least partial completion of the

data access task (Saha, fig. 5, [0060], client buffers of fix size for receiving read data segments or a partial task).

14. For claim 4, Saha-Boyd discloses the invention as in claim 3. Saha-Boyd further discloses:

pre-empting the context resources in the first node allocated for receipt of the first response and for receipt of at least one subsidiary response prior to full completion of the data access task; issuing a second request from the first node by re-activating the context resources that were pre-empted; and (Saha, [0060], client completes reading from a buffer and sends ACK indication that buffer is available for more data);

at a time after issuance of the second request from the first node, allocating context resources to support at least partial completion of the data access task in order to continue repeating issuing at least one subsidiary request and receiving at least one corresponding subsidiary response between the first and second nodes until the data access task is totally complete between the first and second nodes (Saha, [0060], after a partial read completion, send ACK from client to server indicating more data can be sent to the available buffer)

15. For claim 5, Saha-Boyd discloses the invention as in claim 4. Saha-Boyd further discloses allocating context resources to support at least partial completion of the data access task comprises:

reserving the context resources to process subsidiary responses containing data limited in amount according to a preset data allotment identifying an amount of data to be transferred between the first and second nodes (Saha, fig. 5, buffer of fixed size).

16. For claim 6, Saha-Boyd discloses the invention as in claim 1. Saha-Boyd further discloses the second request is at least one of: a different type of data access request than the first request; and received in a request queue that is different that a request queue from which the first request was received (Saha, [0063], [0060], [0061], subsequent queued requests can be read or write).

17. For claim 7, Saha-Boyd discloses the invention as in claim 4. Saha-Boyd further discloses issuing at least one subsidiary request comprises:

calculating a remaining amount of data required to complete the data access task between the first node and the second node (Saha, [0063], queueing remaining data or task segments to be executed, Boyd task consisting of many segments); and

creating at least one subsidiary request to reference at least a portion of the remaining amount of data required to complete the data access task (Saha, [0063], [0060], [0061], subsequent requests amount is the amount of empty buffers)

18. For claim 8, Saha-Boyd discloses the invention as in claim 7. Saha-Boyd further discloses calculating the remaining amount of data comprises:

determining a total completed amount of data processed for the data access task by the first request and associated first response and all subsidiary requests and corresponding subsidiary responses between the first and second node (Saha, [0061], initial total amount of data to be executed in subsequent buffer transfers in a queue);
and

determining the remaining amount of data required to complete the data access task as a difference between an initial amount of data specified by an application request and the total completed amount of data (Saha, [0061], remaining amount of data of a task is the remaining data in the queue and is the difference of initial amount and complete amount).

19. For claim 9, Saha-Boyd discloses the invention as in claim 4. Saha-Boyd further discloses:

the first and second nodes are nodes that utilize channel adapters to exchange the first request and the at least one subsidiary request and the corresponding first response and the at least one subsidiary response (Saha, [0022], Infiniband);

the application request is a remote direct memory access request for the first node to access data in a memory at the second node (Saha, [0022], Remote DMA); and
an initial amount of data specified by the application request is a total amount of data that the first node is to access in the memory at the second node (Saha, a RDMA request inherently has a task size to put in a queue).

20. For claim 10, Saha-Boyd discloses the invention as in claim 9. Saha-Boyd further discloses the first request and the at least one subsidiary request are read remote direct memory access commands issued by the first node to read data in the memory from the second node ([0061], RDMA read).

21. For claim 11, Saha-Boyd discloses the invention as in claim 1. Saha-Boyd further discloses:

establishing a data allotment as a maximum amount of data to be used when responding to requests to transfer portions of data between the nodes, such that if a total amount of data to be transferred between the first node and the second node is greater than the data allotment, the second node provides the first response and the at least one subsidiary response that contain response data that does not exceed the data allotment (Saha, fig. 5, response in fixed size allotments or window); and

allocating context resources in the first node for receipt of the first response and for receipt of the at least one subsidiary response, the context resources allocated to support receipt of responses in an amount that does not exceed the data allotment (Saha, [0061], receive read data in amounts no more than a buffer size).

22. Claims 13-23 are rejected for the same rationale given in claims 1-11 respectively.

23. Claims 12 and 24 are rejected under 35 U.S.C. 103(a) as being unpatentable over Saha-Boyd, further in view of Turner et al. (An Approach For Congestion Control in Infiniband, hereafter Turner).

24. For claim 12, Saha-Boyd discloses the invention as in claim 11. Saha-Boyd does not disclose establishing the data allotment comprises dynamically determining the data allotment between the first and second nodes based on at least one external data allotment event, such that if the at least one external data allotment event occurs, the first and second nodes change a value of the data allotment.

However, Turner discloses the same (fig. 2, section 5.1.3.2, changing a receive window based on a condition)

Therefore, it would have been obvious for one skilled in the art at the time of the invention to combine the teachings of Saha-Boyd and Turner to adjust the amount of maximum data allotment (window control) to, e.g., avoid network congestion (Turner, abstract).

25. Claim 24 is rejected for the same rationale given in claim 12.

Conclusion

26. Any inquiry concerning this communication or earlier communications from the examiner should be directed to Hieu T. Hoang whose telephone number is 571-270-1253. The examiner can normally be reached on Monday-Thursday, 8 a.m.-5 p.m., EST.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Thu Nguyen can be reached on 571-272-6967. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300. Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

/Hieu Hoang/
Examiner, Art Unit 2452